Enabling Network Measurement Portability Through a Hierarchy of Characteristics

(Extended Abstract)

Bruce B. Lowekamp(William & Mary)* Brian Tierney(LBNL) Les Cottrell(SLAC)
Richard Hughes-Jones(Manchester) Thilo Kielmann(Vrije) Martin Swany(UCSB)

December 8, 2002

1 Introduction

Currently, there are many network monitoring tools available to measure a wide variety of network characteristics. For example, there are tools to measure network capacity, available bandwidth, delay, loss, topology, and so on. There are also a number of communities that wish to have a programmatic API for accessing the results of these network measurement tools. For example, network engineers need this measurement data to help identify problems, and so the Internet 2 End-to-End performance initiative [8] was begun to address this need. Another example is the proposed Scalable Internet Measurement Repository [1], a database of measurements, tools, and network experiments.

Network monitoring data is also needed by distributed applications and Grid middleware in order to adapt their behavior based on current network conditions. For example, there may be multiple data servers containing a given file, and the middleware needs to determine which server has the best network connectivity. The combination of many users of measurement data and many ways to obtain that data motivates the development of a common API anyone can use to access the available data.

Building a common API for a wide variety on network monitoring tools requires the ability to classify each tool, and to clearly identify which network characteristics a given tool is measuring. In this paper we define a set of network characteristics that can be measured and a classification hierarchy for these measurements. A classification hierarchy facilitates the determination of common names for these characteristics, and enables the creation of a service that supports general-purpose query protocols, such as SOAP. Such a service provides any program the ability to query for various network monitoring information.

The IETF IPPM working group has defined standard metrics for important networking characteristics [2, 3]. Whereas their work focuses on defining the *right* way to take measurements, our goal is creating a comprehensive system that can be used to categorize *all* measurements that are in use. The network characteristics hierarchy presented here is a first step towards addressing the issue of providing a common framework to name and query network measurements.

This extended abstract presents a condensed version of work done by members of the "Network Measurements Working Group," which is part of the Global Grid Forum. The full version of the paper will provide more details of the classification system described in the first few sections, and include an expanded evaluation section describing our current projects making use of the hierarchy and discussing how our clas-

*Corresponding author: lowekamp@cs.wm.edu

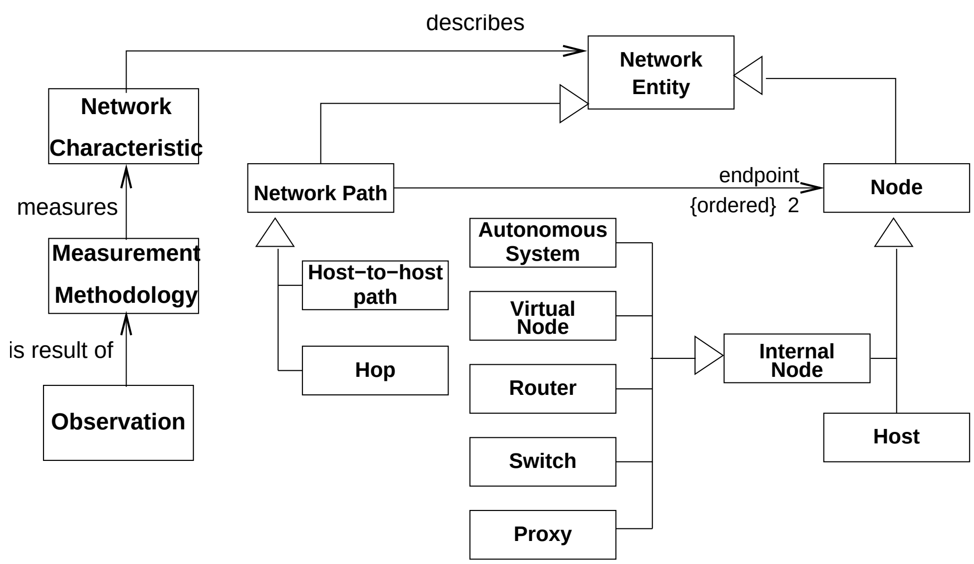

Figure 1: The relationship between observations, measurement methodologies, characteristics, and entities. Network entities are divided into nodes and paths. A node can be one of several components, ranging from an autonomous system (AS) to a host or switch. A path is a unidirectional connection from one node to another node, indicated by the ordered pair of endpoints in the diagram. Paths and nodes are both annotated with protocol and QoS attributes so that network behavior with different types of traffic can be represented. In our UML notation, the boxed arrows indicate an inheritance, or an "is-a" relationship, and the open arrows indicate an attribute.

sification system can be used with a variety of current monitoring systems, such as NWS, Remos, PingER, and TopoMon.

2 Terminology

We define three terms, ranging from general to specific. *Network characteristics* are the intrinsic properties of a portion of the network that are related to the performance and reliability of the network, e.g. hop capacity. *Measurement methodologies* are the means and methods of measuring those characteristics, e.g. as implemented by tools like pathchar. An *observation* is an instance of the information obtained by applying the measurement methodology.

Figure 1 illustrates the relationship between these terms. While the difference between a network entity and a characteristic of that entity is clear, differentiating between characteristics and measurement methodologies is frequently more difficult. Note that a characteristic is not necessarily associated with a single number. For instance, packet loss is an important characteristic of paths. However, loss can be expressed generally as a fraction of all traffic sent, or as a loss-pattern with detailed statistical properties. The intuitive sense is that all measurement methodologies under a particular characteristic should be measuring "the same thing," and could be used almost interchangeably, while measurements methodologies under separate characteristics are not directly interchangeable. (Although they may be used along with other characteristics to derive those values.)

To determine if a particular concept is a characteristic or a measurement methodology, the most important factor is whether the technique used to make the observation influences the value measured. If there are different methods to observe identical or similar concepts that result in different values, then the concept may be a characteristic, but the techniques are measurement methodologies.

Measurement methodologies and observations retain their definitions described by IPPM. In some sense, the IPPM uses *metric* where we use characteristic, however we found that the multiple meanings of *metric* introduced too much confusion, therefore we use *characteristic* and define it more precisely for our purposes.

3 Measurements Representation

There are two components necessary to describe a network measurement. The first is the characteristic being measured. The second is the network entity that the measurement describes—the path, hop, host, router, etc.

Developing a way for measurement systems to interchange information about network measurements requires a uniform way of representing the measurement, the network entity it measures, and the conditions under which the measurement is performed. Our representation combines the network entity with the conditions under which measurements are taken, such as protocol and QoS. This allows us to represent all possible measurements. Note that entities are defined by the measurement systems reporting measurements for them—there may be duplicates and, in particular for topologies, overlapping entities reported by different systems.

3.1 Network Entities and Topology

As networks are best represented in graph form, network entities are divided into nodes and paths, as illustrated in Figure 1. A node does not necessarily correspond to a single physical entity, but can also represent an autonomous system or a virtual node. A path exists between the two nodes used as measurement endpoints.

Nodes and paths are annotated with the attributes used in taking the measurement, including protocol stack (e.g., TCP over IPv4) and QoS level. Note that the attributes are not merely descriptive of the node or path, but form a tuple that defines a unique node or path. In other words, the entity model superimposes multiple topologies over the actual physical nodes and paths, allowing the behavior of the network under different types of traffic to be characterized separately. Also note that some characteristics, such as capacity, may be most appropriately described without specifying either a protocol or QoS level.

Paths are used to represent the connection between any two nodes in the network. A path can be anything from an end-host to end-host connection across the Internet to a single link between two Ethernet switches. Representing a network's topology with link-level paths allows us to build a detailed model of a network's structure. The combination of types of nodes and paths can be used to represent both physical topologies [4, 7, 10, 11] and functional topologies [5, 9, 12, 13]. The power of our representation is its ability to support diverse topologies using a single system.

A path may be either a host-to-host path or a single "hop." A measurement system reporting only characteristics of end-to-end measurements would use only host-to-host paths, whereas a system attempting to provide a topology representation would use hops to build that topology. There is no difference in terms of the attributes or characteristics applied to each type of path.

The hoplist characteristic (see Figure 2) allows a path, such as an end-to-end path, to be subdivided into the hops that form the path. For example, a path between two end-hosts could have a hoplist containing the set of hops between each of the routers along the end-to-end path. Each member of the hoplist is itself a hop, which is also a type of path as shown in Figure 1. End-to-end paths are not the only paths that can

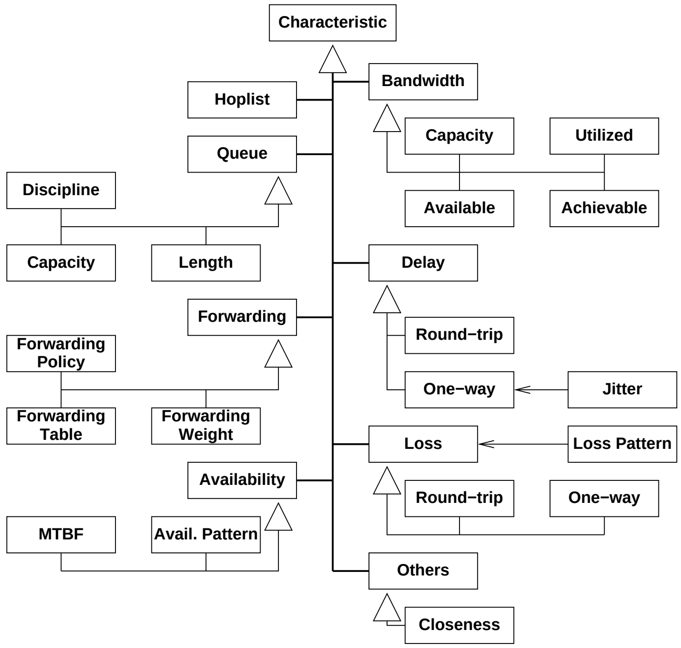

Figure 2: Network characteristics that can be used to describe the behavior of network entities.

be divided into hops with a hoplist. For example, one of the above router-to-router hops could be further subdivided into hops representing the links between Ethernet switches in a bridged LAN.

3.2 Characteristics

The characteristics hierarchy is shown in Figure 2. Any number of these characteristics may be applied to a network entity. Note that some characteristics, such as route or queue information, are only sensible when applied to either paths or nodes.

The full version of this paper contains detailed definitions for all of the boxes in Figure 2, and includes a description of the issues involved in actually measuring each characteristic. Due to space constraints, we only include here very condensed versions of the characteristics for bandwidth, delay, and loss.

Bandwidth: Bandwidth is defined most generally as data per unit time. However, the "bandwidth of a path" is not a precisely defined term and specific characteristics must be clearly defined prior to discussing how to measure bandwidth.

There are four characteristics that describe bandwidth:

Capacity: Capacity is the theoretical maximum link-layer throughput that the path has available, when there is no competing traffic.

Utilization: The aggregate of all traffic currently on that path.

Available Bandwidth: The current unutilized bandwidth. This can be measured directly or estimated from capacity and utilization.

Achievable Bandwidth: The throughput that an application might achieve on a given host type, OS type, protocol configuration (e.g., TCP buffers), and available bandwidth on a path. This is what tools such as iperf provide.

Delay: Delay is the time for a packet to travel between two nodes, and may be divided into one-way and round-trip delay. The variance in one-way delay is commonly referred to as jitter.

Loss: Loss also may be one-way or round-trip. A loss pattern can be used to represents more detailed observations than a simple fraction.

4 Evaluation

The characteristics hierarchy we have described is being evaluated through use in network measurement archiving systems being studied at LBNL and UCSB. We have also ensured that the characteristics hierarchy is capable of representing the measurements taken by a wide variety of systems we have designed and implemented, encompassing a range of systems from host-to-host measurements to topology-based measurements.

Based on the hierarchy of characteristics described in this paper, we have implemented a SQL-based network measurements archive. The following table shows a subset of the measurements that are stored in this archive. The existence of this characteristics hierarchy made it much easier to determine monitoring event names and suggested the structure of the tables. All monitoring events in the table can be queried by specifying a source and a destination, where this source-destination pair may be end hosts or any network entity. The naming convention used for the network monitoring events are based on work by the Global Grid Forum DAMED working group [6], which considered an earlier version of our characteristics hierarchy in developing their events.

Monitoring Event Name	**Units**	**Attributes**
bandwidth.achievable.tcp.singleStream	Mbits/sec	test duration, measurement tool, TCP buffer size, etc.
bandwidth.achievable.tcp.multiStream	Mbits/sec	same
bandwidth.achievable.udp.singleStream	Mbits/sec	same
delay.roundTrip	ms	protocol used (e.g.: ICMP), packet size
loss.oneWay	percentage	measurement tool. time interval
hop.bandwidth.capacity	Mbits/sec	measurement tool
hop.bandwidth.utilized	Mbits/sec	measurement tool, time interval
hoplist	list of IPs	measurement tool

A clear example of using the measurement hierarchy involves queries posted against these databases. For instance, to check for the availability of measurements of achievable bandwidth, the search might match any event name beginning with “bandwidth.achievable.tcp.*”.

Another qualitative effect of this naming scheme is the ability to share measurement data across archive systems. The use of a a consistent monitoring event naming standard across monitoring data repositories facilitates data sharing. Coupled with metadata about measurement parameters (duration, packet size, etc.) these measurements are unambiguously identified and can be used by others. Many organizations are collecting data based on the conventions outlined. The final paper will present interoperability results as well.

The full paper will have a detailed evaluation section including further implementation experience.

5 Conclusion

Without a methodology for classifying different types of network measurements, it is very difficult to construct databases and APIs to provide access to network monitoring data. The portability of grid applications and the power of shared measurement repositories both depend on the development of systems and techniques for supporting true portability of network measurements. The network characteristics hierarchy presented here is a first, essential step towards addressing this issue.

References

[1] Mark Allman, Ethan Blanton, and Wesley M. Eddy. A scalable system for sharing internet measurements. In *Proceedings of the Passive and Active Measurements Workshop 2002 (PAM2002)*, March 2002.

[2] G Almes, S Kalidindi, and M Zekauskas. A one-way delay metric for IPPM. RFC2679, September 1999.

[3] G Almes, S Kalidindi, and M Zekauskas. A one-way packet loss metric for IPPM. RFC2680, September 1999.

[4] Yuri Breitbart, Minos Garofalakis, Cliff Martin, Rajeev Rastogi, S. Seshadri, and Avi Silberschatz. Topology discovery in heterogeneous IP networks. In *Proceedings of INFOCOM 2000*, March 2000.

[5] Mark Coates, A O Hero III, Robert Nowak, and Bin Yu. Internet tomography. *IEEE Signal Processing Magazine*, 19(3):47–65, May 2002.

[6] Discovery and Monitoring Event Description (DAMED) Working Group, Global Grid Forum. `http://www-didc.lbl.gov/damed/main.html`.

[7] Ramesh Govindan and Hongsuda Tangmunarunkit. Heuristics for internet map discovery. In *IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000.

[8] Internet2 end-to-end performance initiative. http://www.internet2.edu/e2eperf/.

[9] Sugih Jamin, Cheng Jin, Yixin Jin, Danny Raz, Yuval Shavitt, and Lixia Zhang. On the placement of internet instrumentation. In *IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000.

[10] Bruce Lowekamp, David R. O'Hallaron, and Thomas Gross. Topology discovery for large ethernet networks. In *Proceedings of SIGCOMM 2001*. ACM, August 2001.

[11] Venkata N. Padmanabhan and Lakshminarayanan Subramanian. An investigation of geographic mapping techniques for internet hosts. In *Proceedings of ACM SIGCOMM 2001*, pages 173–185, 2001.

[12] Gary Shao, Fran Berman, and Rich Wolski. Using effective network views to promote distributed application performance. In *Proceedings of the 1999 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'99)*, 1999.

[13] Wolfgang Theilmann and Kurt Rothermel. Dynamic distance maps of the internet. In *IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000.